Soc Choice Welfare (1993) 10: 161-175

Social Choice and Welfare

Almost all social choice rules are highly manipulable, but a few aren't

Jerry S. Kelly*

Syracuse University, Syracuse, NY 13244-1090, USA

Received November 5, 1990 / Accepted September 17, 1992

Abstract. Explores, for several classes of social choice rules, the distribution of the number of profiles at which a rule can be strategically manipulated. In this paper, we will do *comparative* social choice, looking for information about how social choice rules compare in their vulnerability to strategic misrepresentation of preferences.

1. Introduction

We must start by sorting out some conceptual problems. How shall we measure the extent to which different social choice rules can be manipulated by voters submitting ballots that differ from their true preference rankings? One possible answer would be to add by stages new constraints that must be satisfied before manipulation can take place. The degree of manipulability of a rule would then be measured by the stage in this process at which it becomes *strategy-proof*, where **no** manipulations can take place.

An example of this approach is found in the study of counterthreats (Pattanaik 1978). Suppose that we agree manipulations will not take place if the manipulator can be dissuaded by a credible counterthreat. If a second individual could respond to the initial manipulation by a counter-manipulation, forcing an alternative that the original manipulator likes less than what would have occured without his action, then the original manipulation would be deterred. A rule is *strategy-proof with counterthreats* if, at every profile where the rule is manipulable, then for every manipulation there is a credible counter-manipulation. If two social choice rules, f_1 and f_2, are both manipulable, so that neither is strategy-proof, but f_2

* I am indebted to Universidad Internacional Menendez Pelayo for an invitation to present an early version of this paper at a June, 1989 conference at Valencia, Spain. Thanks for comments go to the participants of that conference, especially Salvador Barbera; also for comments at seminars at Syracuse University and the University of Rochester.

is strategy-proof with counter-threats, we could say that f_1 is more manipulable than f_2.

A first difficulty with this approach, is that we have just a single weakening, rather than a lengthy sequence of weakenings. As a result, we obtain only a very coarse partition of social choice rules into a most two degrees of manipulability. That situation, of course, might be improved with introducing additional weakenings of strategy-proofness. More serious is that the partitioning may well be trivial; no one has yet presented a social choice rule that is strategy-proof with counterthreats but not strategy-proof. We don't have any non-trivial cases where f_1 is more manipulable than f_2.

An alternative weakening of strategy-proofness (Kelly 1988) by imposing constraints on manipulability starts with a Kemeny (1959) distance function defined on the set of preference orderings. We imagine that other voters have enough experience with you to know if you are submitting a preference very far, a distance d or greater, from your true ordering. So you can only get away with "local" manipulations to orderings less than d away from your true ordering. For a social choice rule, f, let $D(f)$ be the shortest distance between a true preference and its misrepresentation. Then we will say rule f_1 is less manipulable than f_2 if $D(f_1) > D(f_2)$, i.e., if manipulations at f_2 are less detectable than those at f_1. Clearly, there exists here the logical possibility of more than the two degrees of manipulability that were possible in the classification by counterthreats. However, that turns out to be quite deceptive; proofs of the Gibbard-Satterthwaite theorem show that manipulable rules can always be manipulated by misrepresentation to an ordering that is the logically minimal distance[1] from a true ordering. There is only one degree.

Each of the proposals in this first part results in a failure to spread rules out in a variety of degrees of manipulability. A second approach is to say that rule f_1 is less manipulable than f_2 if manipulation of f_1 requires voters to have more information than manipulation of f_2. One suggestion along these lines has come from Nurmi (1987), who also presents his orderings of some standard rules. Unfortunately, so far no one has given precise quantitative definitions of a "measure of the information" that individual i needs in order to manipulate at profile u in the domain of rule f. And no one has suggested how to aggregate such a measure over all profiles and all individuals for f. Thus it is not yet possible to prove that any presented orderings are correct; nor is it possible to calculate the minimum possible degree of manipulability in a class of rules or do anything else in the way of rigorous theory.

There is, however, a third approach to a measure of degree of manipulability that both spreads rules out in a variety of degrees and provides a basis for detailed calculation and rigorous theory development. If we fix the number of individuals and the number of alternatives and specify whether weak orders (with non-trivial indifference) as well as strong preference orderings are allowed, and if we specify that social choice rules must satisfy *universal domain* (the rule's domain consists of all logically possible profiles of preferences), then we say rule f_1 is more manipulable than rule f_2 if there are more profiles at which f_1 is manipulable. (Notice that a profile is counted just once even if there are several individuals who would manipulate there.)

[1] The minimal Kemeny distance between two strong orders is 1; between two weak orders is 1/2.

	xyz	xzy	yxz	yzx	zxy	zyx
xyz	x	x	ⓧ	ⓨ	x	ⓧ
xzy	x	x	x	ⓧ	ⓧ	ⓩ
yxz	ⓧ	x	y	y	x	y
yzx	ⓨ	ⓧ	y	y	z	ⓨ
zxy	x	ⓧ	x	z	z	z
zyx	ⓧ	ⓩ	y	ⓨ	z	z

Fig. 1

2. Comparisons in class of unconstrained rules

Very little is known about how rules compare on this profile count measure of manipulability. Gibbard (1973) and Satterthwaite (1975) have shown us that for most interesting classes of rules the degree of manipulability is at least one; there exists at least one profile at which each rule can be manipulated by at least one voter. Blair (1981) has determined information useful for establishing lower bounds in a narrow class of rules. Lepelley and Mbih (1987) calculate, for the plurality rule, the number of profiles at which that rule can be manipulated by coalitions (rather than, as here, by individuals). That is almost all we know rigorously. There is some tradition that manipulation is common (Riker 1982) but that certain rules, like the Borda rule, are especially manipulable (see Chamberlin 1985; Nitzan 1985 and Saari 1990).

In an oft-told story (see Black 1968; Lacroix 1800 and Mascart 1919) about Borda's invention of his voting procedure, he is confronted with the criticism that voters would find it advantageous to submit ballots in which they falsely place – at the bottom – the strongest opponents to their favorite candidates. Borda replies: "My scheme is only intended for honest men." Living now in the light of the Gibbard and Satterthwaite theorem, we know that a large and important class of social choice rules will share with Borda's the possibility of at least one situation where strategie manipulation is both possible and advantageous. But many commentators treat the Borda procedure as if it were especially vulnerable to false ballot submission. One purpose of this paper is to improve our intuitions about the comparative degree of manipulability of social choice procedures.

Consider Fig. 1 above, which illustrates a version of the Borda rule[2]. Here there are three alternatives: x, y and z. There are two voters, each with strong preferences; #1's possible ballots are listed down the left side while #2's are listed across the top. We focus on *resolute* rules, i.e., social choice rules where at each profile of submitted ballots, at each box in the diagram, exactly **one** candidate is chosen. Put differently, we are assuming that the tie breaking method is built into our social choice rule. In our Fig. 1 example, we have supplemented the Borda procedure, which may generate ties, with the rule that ties will be broken alphabetically: if x and z are both Borda winners, x will be chosen. Of the 36 profiles, 14 (those with circles) have outcomes at which at least one voter would have an incentive to submit a false ballot.

[2] Diagrams of this kind were introduced by Feldman (1980).

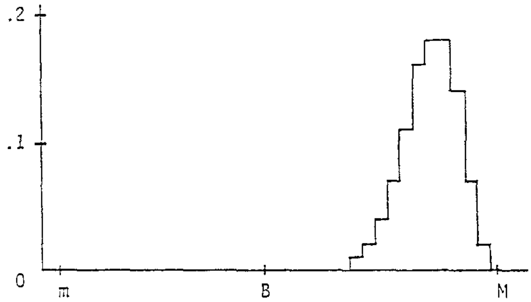

Fig. 2. Degrees of manipulability: Two agents, three alternatives, strong orders

How bad is this performance by the Borda count rule as compared with other rules that work with two individuals having strong orders on three candidates? Let's start, very naively, by looking at the full class of social choice rules indicated by the Gibbard-Satterthwaite theorem. This is the class of rules that

i) are resolute;
ii) have no dictator;
iii) satisfy universal domain; and
iv) have at least three alternatives in their range.

Throughout the rest of this paper we will assume that these four conditions are always satisfied.

Since each of the 36 boxes can be filled in three ways, there are 3^{36} ways of filling all boxes. Of these, 2^{36} are filled with just x and y - no occurrences of z; a similar number are filled with just x and z; similarly with y and z. But there is a mild double counting here; the rule that always gives x will be among the 2^{36} filled with just x and y and also among those filled with just x and z. Hence there are $3^{36}-3 \cdot 2^{36}+3$ ways of filling the boxes with all three of x, y and z. Of these, two ways are dictatorial. Hence the class of non-dictatorial rules with three alternatives in the range has $3^{36}-3 \cdot 2^{36}+1 \approx 1.5 \times 10^{17}$ members. This number is large enough to discourage us from any idea of an exhaustive rule-by-rule determination of the population distribution of degrees of manipulability. For this reason, much of what is reported here is the result of sampling experiments.

Before we turn to those experiments, however, it will help us to learn a little more about the possible extent of manipulability of rules in this domain. The maximum degree of manipulability in the class, the largest number, M, of profiles at which at least one voter has an incentive to manipulate, is 36. Just think about the rule that always chooses #1's least preferred alternative. The smallest number, m, of profiles at which at least one voter has an incentive to manipulate, is at least 1 - that is the import of the Gibbard-Satterthwaite theorem. An analysis (Kelly 1988) shows that $m=2$.

In the first sampling experiment we analyze, a rule is constructed by independently filling each box with x, y and z (each with probability 1/3) and then discarding rules with dictators or small ranges. In this fashion, 50 000 rules were constructed and each was examined to calculate the number of profiles at which they were manipulable. The resulting frequency distribution is shown in Fig. 2.

This distribution is bunched up close to M; the mean number of profiles at which the rule is manipulable is 30.855. Most social choice rules are highly manipulable. Certainly the Borda rule, with its 14 profiles at which the rule is manipulable – indicated on the diagram with a B – does well relative to this whole class. Of the 50 000 rules, only one did as well, with 14 profiles at which the rule is manipulable.

The bunching up of sampled degrees near the maximum, M, will be one of the two main pieces of the story we will be telling about distributions. Perhaps the most dramatic piece of the story. But the other half may ultimately be more important. Although rules with lower degrees of manipulability, with degrees near m, may be much rarer, though they may be distributed "thinly upon the ground", they **are** there. We can illuminate this kind of result by pointing out that a proof (Kelly 1989) shows that for every integer i between m and M there is a rule in this class with exactly i profiles at which the rule is manipulable. This is the property of **interjacency** for this class.

Of course, the observation that the Borda procedure does well within this class is not very heartening since nearly all rules in this class are terrible anyway. Most rules, in addition to being highly manipulable, will fail to have many other desirable properties. This suggests altering our experiment to restrict attention to just those rules satisfying some list of desirable criteria.

Before we do this, however, let's pursue the results of the unconstrained case a little further to see that they are robust under several modifications. In particular, we shall want to change the number of voters, change the number of candidates and work with weak orders as well as strong.

A. Weak orders

Let us first retain three alternatives and two voters. The only change is that we will allow voters to have weak orders over alternatives – they can express indifference on their ballots. For this class of social choice rules, each with a domain of 169 profiles, analysis shows that m is still 2 and $M = 168$ (manipulation is not possible at the one profile where both voters are indifferent among all three alternatives). To see where Borda's rule fits in this distribution, we must first introduce a modification to take care of profiles of weak preferences since Borda's procedure was originally described only for the case of strong orders. A modification of this sort was done by Black (1959) and by Luce and Raiffa (1957). Associated with each individual i and each alternative x, there is a *Borda score*

$$B(i,x) = \#\{y/xP_iy\} - \#\{y/yP_ix\} \ .$$

The Borda rank of alternative x is then obtained by adding up Borda scores over individuals:

$$r(x) = \sum_i B(i,x) \ .$$

Finally, the Borda rule selects those alternatives with highest Borda rank. This modified Borda rule with ties broken alphabetically is manipulable at exactly 42 profiles.

In our sampling experiment, 50 000 rules in this class were obtained and the resulting distribution is pictured in Fig. 3. The sample range was from 113 to 149, so that all degrees were much higher than the degree of the Borda rule. The sample mean of degrees was 132.58; the distribution is bunched up near M.

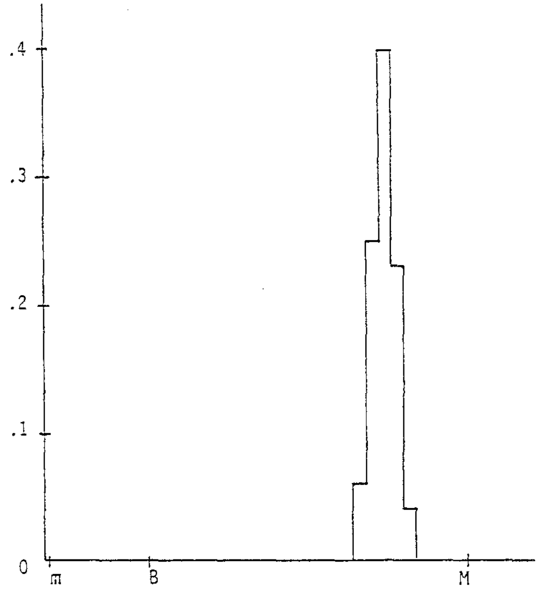

Fig. 3. Degrees of manipulability: Two agents, three alternatives, weak orders

Again, the second half of the story is an interjacency result. While most rules in this class are very manipulable, there are some rules at **each** lower degree of manipulability. For this class of social choice rules, with $m=2$, $M=168$ we discover that for every i between m and M there is a rule in this class with exactly i profiles at which the rule is manipulable.

B. Three individuals

Let us next increase the number of voters to three while still retaining the assumptions that there are three candidates over which each voter has a strong preference ordering. The Borda rule, with ties broken alphabetically, is manipulable at 51 profiles. M, the maximum number of profiles at which a rule in this class is manipulable is 216 (think about the rule that always chooses # 1's worst). The minimum, m, is not known for sure. The rule from this class with the smallest known number of profiles at which the rule is manipulable has five such profiles. This rule is shown in Fig. 4. It is illustrated to show the very poor quality of rules that minimize manipulability.

The undesirability of manipulation-minimizing rules is the reason for the importance of interjacency. Starting from a manipulation-minimizing rule, we would like to trade a little more manipulability in exchange for improvement in desirability on other criteria. But if there a big gaps in degrees of manipulability, we will have to trade only large increases in manipulability for other improvements. For the class we consider here, for every integer i between 5 and 216, there is a rule that has exactly i profiles at which it is manipulable. A result of this sort will be called *conditional interjacency* since it is only interjacency conditional upon 5 being the minimum possible number.

	x	x	y	y	z	z
	y	z	x	z	x	y
	z	y	z	x	y	x
xyz						
xyz	x	x	x	x	x	x
xzy	x	x	x	x	x	x
yxz	x	x	x	x	x	x
yzx	x	x	x	x	x	x
zxy	x	x	x	x	x	x
zyx	x	x	x	x	x	x
xzy						
xyz	x	x	x	x	x	x
xzy	x	x	x	x	x	x
yxz	x	x	x	x	x	x
yzx	x	x	x	x	x	x
zxy	x	x	x	x	x	x
zyx	x	x	x	x	x	x

	x	x	y	y	z	z
	y	z	x	z	x	y
	z	y	z	x	y	x
yxz						
xyz	y	y	y	y	y	y
xzy	y	y	y	y	y	y
yxz	y	y	y	y	y	y
yzx	y	y	y	y	y	y
zxy	y	y	y	y	y	y
zyx	y	y	y	y	y	y
yzx						
xyz	y	y	y	y	y	y
xzy	y	y	y	y	y	y
yxz	y	y	y	y	y	y
yzx	y	y	y	y	y	y
zxy	y	y	y	y	y	y
zyx	y	y	y	y	y	y

	x	x	y	y	z	z
	y	z	x	z	x	y
	z	y	z	x	y	x
zxy						
xyz	z	z	z	z	z	z
xzy	z	z	z	z	z	z
yxz	z	z	z	z	z	z
yzx	z	z	z	z	z	z
zxy	z	z	z	z	z	z
zyx	z	z	z	z	z	z
zyx						
xyz	z	z	z	(z)	z	z
xzy	z	z	z	z	z	z
yxz	z	z	z	(z)	z	z
yzx	(z)	z	(z)	(y)	z	z
zxy	z	z	z	z	z	z
zyx	z	z	z	z	z	z

Fig. 4

For three voters with strong orders on three candidates, 50 000 independently generated rules were obtained and a count made of the number of profiles at which they are manipulable. The results are given in Fig. 5. Clearly, the distribution is bunched up near M; the sample mean is 204.28. By comparison, the value of 51 for the Borda rule seems quite low. None of the sampled rules had degree near that of the Borda rule; the lowest was 184.

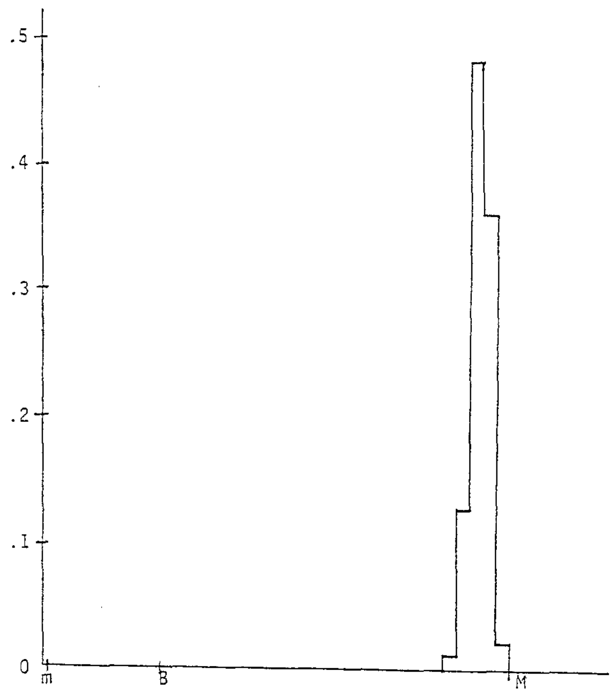

Fig. 5. Degrees of manipulability: Three agents, three alternatives, strong orders

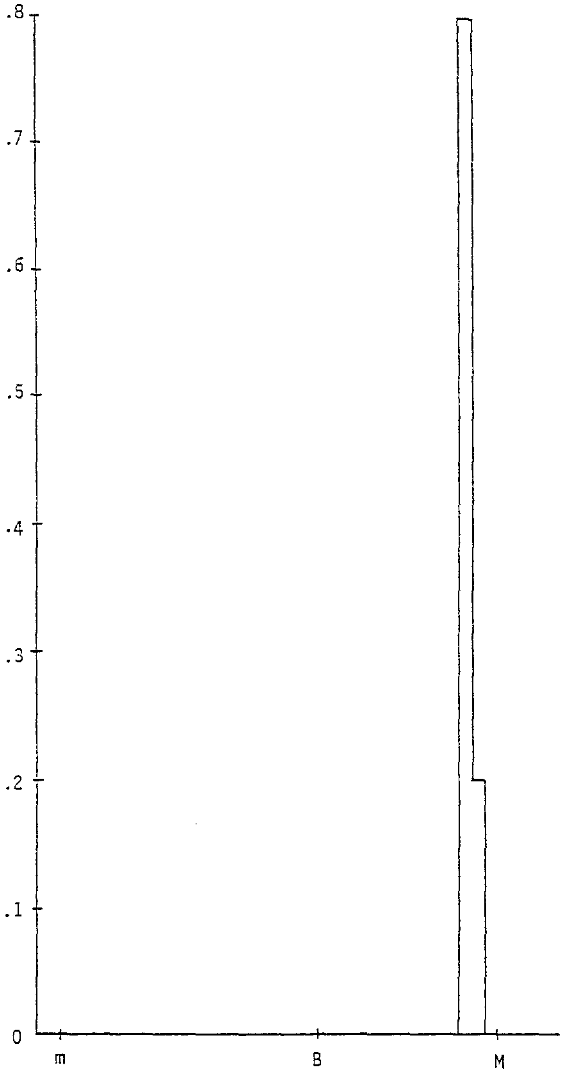

Fig. 6. Degrees of manipulability: Two agents, four alternatives, strong orders

C. Four alternatives

As a last test of robustness, let us retain two individuals with strong orders but we increase the number of alternatives to four, of which at least three must be in the range. The two individual, four alternative Borda rule with ties broken alphabetically is manipulable at exactly 348 of the 576 profiles. For this class, $M = 576$ while m is unknown; the smallest known degree of manipulability is 12.

In the corresponding sampling experiment, 50 000 rules in this class were obtained and we have the usual theme reflected sharply in Fig. 6: degrees are bunched up near M. All sampled degrees were well above the degree of the Borda rule. Also for this class, conditional interjacency holds; for every i between 12 and 576 there is a rule manipulable at exactly i profiles.

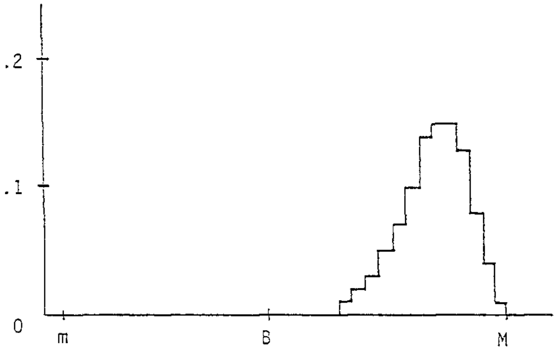

Fig. 7. Degrees of manipulability: Two agents, three alternatives, strong orders, anonymous

3. Symmetry constraints

As we observed earlier, we learn only a little bit by seeing how these degrees of manipulability have been bunched up near M because these classes are unconstrained and most rules in each class, in addition to being highly manipulable, will fail to have many other desirable properties. So we now consider looking at classes of rules satisfying other constraints besides resoluteness, universal domain, range size at least three, and non-dictatorship.

A distinction should be drawn between constraints known to be integrally connected with manipulability and those that are more independent. For example, Muller and Satterthwaite (1977) have shown the equivalence of strategy-proofness with a strong positive responsiveness condition. This suggests that responsiveness constraints, and constraints like the Pareto condition that are closely related to responsiveness conditions, are likely to push the degree distribution from M down toward m. By way of contrast, constraints for which there are no theorems connecting them with strategy-proofness might be expected to leave some characteristics of the distributions unaffected.

A. Anonymity

To illustrate this latter point, suppose that we focus on the subclass of rules that satisfy a symmetry property, anonymity[4], which ensures a kind of equal treatment of voters. The minimum number of profiles, m, at which the rule is manipulable remains 2 and the maximum, M, is 36. Interjacency also holds for this class; for all integers i between m and M, there is an (anonymous) social choice rule in this class manipulable at exactly i profiles.

A set of 50 000 rules satisfying all the conditions was generated; Fig. 7 shows the relative frequency of the degree of manipulability. The manipulability degrees remain bunched up near M; most anonymous social choice rules are very manipulable; the sample mean for this class is 30.616, not much less than the 30.855 mean in the unconstrained case. There is, however, a clear **spreading out** of the

[4] Let σ be any permutation on individuals. Then if $u=(p_1,p_2,\ldots,p_N)$ we define $\sigma(u)=(p_{\sigma(1)},p_{\sigma(2)},\ldots,p_{\sigma(N)})$. A SCR satisfies *anonymity* if at all profiles u and all permutations σ, $f(\sigma(u))=f(u)$.

	xyz	xzy	yxz	yzx	zxy	zyx
xyz	x	x	(x)	(y)	x	(x)
xzy	x	x	x	(x)	(x)	(z)
yxz	(y)	(x)	y	y	(y)	y
yzx	y	(y)	y	y	(z)	(y)
zxy	(x)	(z)	(z)	z	z	z
zyx	(z)	z	(y)	(z)	z	z

Fig. 8

distribution about the sample mean and 15 of the 50 000 rules had degree less than that of the Borda procedure.

B. Neutrality

A related symmetry condition is neutrality[5], which ensures a kind of equal treatment of candidates. Several new issues arise when we focus on the case with two individuals, three alternatives and satisfaction of neutrality. First of all, the Borda rule in Fig. 1 is not neutral because the tie-breaker is not neutral, depending on the alphabetical order of the alternatives – x has a privileged position. So let us construct a new rule, based on the usual Borda procedure – which generates ties – but use #1's order as the tie breaker. We gain neutrality at the expense of anonymity. The rule constructed in this fashion is shown in Fig. 8 and is manipulable at 18 profiles.

The next issue is that this class of neutral rules is small enough that we can forego sampling in favor of examining the exact distribution. There are 727 neutral rules with no dictator. Within this class, $m = 6$ and $M = 36$. The (population) mean is 30.95, still quite close to the (sample) mean of 30.855 in the unconstrained class. Only four of the 727 rules have lower degree than the Borda rule of Fig. 8.

The simplest kind of interjacency will now fail. If f is manipulable at profile u, then f is also manipulable at any profile obtained from u by a permutation of the alternatives. Since there are six such permutations, any degree of manipulability must be a multiple of six. The exact distribution is displayed in Fig. 9 where we again see a distribution bunched up close to M. While simple interjacency fails, we still have a kind of **group interjacency;** for every multiple $6 \cdot i$ from $m = 6$ to $M = 6 \cdot 6$, there is at least one rule with degree $6 \cdot i$.

C. Anonymity and neutrality

As a last example of symmetry constraints, we now examine the class of rules with three alternatives and two individuals, that satisfy both anonymity and neutrality. Of course, as related above, degrees of manipulability must be mul-

[5] Let θ be any permutation on X. Then if $p \in \mathbf{P}$, we define $\theta(p)$ by the following rule for all $a, b \in X$: apb if and only if $\theta(a)\,\theta(p)\,\theta(b)$. Then, if u is the profile $(p_1, p_2, \dots, p_N)$ we define $\theta(u)$ to be the profile $(\theta(p_1), \theta(p_2), \dots, \theta(p_N))$. Finally, a SCR f satisfies *neutrality* if at all profiles u and all permutations θ, $f(\theta(u)) = \theta(f(u))$.

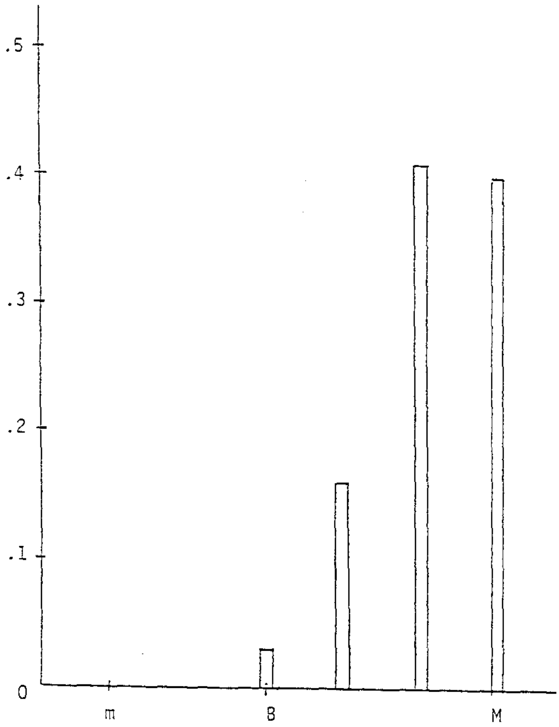

Fig. 9. Degress of manipulability: Two agents, three alternatives, strong orders, neutral

tiples of 6 and can be exactly counted. There are three such rules with multiplicity $m = 24$ and six rules with multiplicity $M = 30$, for a mean of 28. In some trivial sense, group interjacency holds.

4. Pareto

As observed above, Muller and Satterthwaite have shown that strategy-proofness is equivalent to a strong responsiveness condition. This leads us to expect that imposing a responsiveness condition or a related constraint on a social choice rule will have significant impact on degree distribution, probably pushing degrees closer to m.

To illustrate this, we first work with a condition closely related to responsiveness, the *Pareto* condition: If there is some alternative that both individuals prefer to candidate A, then A is not chosen. We return to the case of two individuals, three alternatives, strong orders and no symmetry conditions, but where now the rules must all satisfy the Pareto condition. Although the minimum number, $m = 2$, of profiles at which the rule is manipulable remains the same as in the unconstrained case, the maximum falls to 24; at the 12 profiles where both ballots have the same top candidate, that candidate must be chosen and the rule is not manipulable there.

A set of 50 000 rules satisfying the Pareto condition was generated and each was examined to calculate the number of profiles at which it was manipulable.

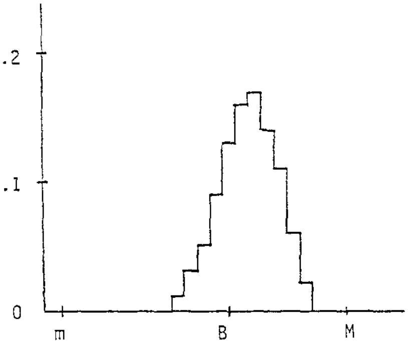

Fig. 10. Degrees of manipulability: Two agents, three alternatives, strong orders, Pareto

The resulting frequency distribution is shown as Fig. 10. The mean has shifted from 30.855 to 16.496, closer to the degree, 14, of the Borda rule, which is now more nearly typical.

Interjacency still holds; for every integer i between 2 and 24 there is a rule satisfying the Pareto condition with exactly i profiles at which the rule is manipulable. This result is fairly robust. If we increase the number of individuals to three, M changes to 192 (with a domain of 216); the smallest known degree is 5. Conditional interjacency holds; for every integer i between 5 and 192 there is a rule satisfying the Pareto condition with exactly i profiles at which the rule is manipulable. If we return to two individuals but allow weak preferences as well as strong, $M=72$ (of 169 profiles); the smallest known degree is 3. For nearly every integer i between 3 and 192 there is a rule satisfying the Pareto condition with exactly i profiles at which the rule is manipulable; the only exception is $i=4$. It is not known if there is a rule on weak preferences satisfying the Pareto condition with excactly 4 profiles at which the rule is manipulable.

A. Anonymity

Suppose that now we consider what happens if we combine the Pareto condition with symmetry conditions. To start, we examine the class of rules with two individuals, three alternatives, strong orders, satisfying both anonymity and the Pareto-condition. If $u=(P_1, P_2)$ is a profile at which f is manipulable, the two individuals could not have the same ordering at u, i.e., $P_1 \neq P_2$, otherwise the Pareto condition would be violated. But then the profile $u' = (P_2, P_1) \neq u$ is also manipulable, by anonymity. Hence the number of profiles at which the rule is manipulable must be a multiple of 2; finally, $m=2$ and $M=24$.

The distribution of degrees from a sample of 50 000 is presented in Fig. 11. The mean of this sample is 16.48 which is very close to the mean of the previous experiment, 16.496. Adding the anonymity condition to Pareto doesn't have much impact. But this mean is quite a bit smaller than the mean of 30.616 that we obtained when we required anonymity but not the Pareto condtion. The Pareto condition has pushed the distribution away from M.

For this class, we also have satisfaction of group interjacency; for every $2 \cdot i$ between 2 and 24, there is an anonymous rule satisfying the Pareto condition with exactly $2 \cdot i$ profiles at which the rule is manipulable. Regarding robustness

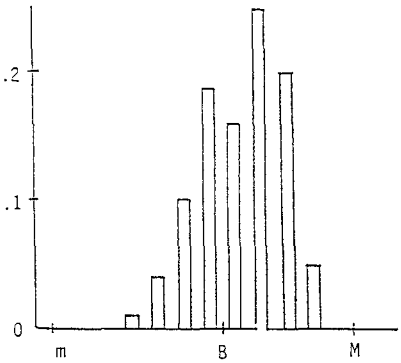

Fig. 11. Degrees of manipulability: Two agents, three alternatives, strong orders, anonymous and Pareto

of this interjacency claim, we may consider relaxing the domain to allow weak orders as well as strong. The domain now has 139 profiles; analysis shows that $M=72$. The least manipulable rule found so far in this class has 10 profiles at which the rule is manipulablemanipulable. We almost have conditional group injacency: for nearly every $2 \cdot i$ between 10 and 72 there is a rule in this class with exactly $2 \cdot i$ profiles at which the rule is manipulable. The only gap is that there is no known rule in this class with exactly 70 profiles at which the rule is manipulable.

B. Neutrality

Similarly, we examine the class of rules with two individuals, three alternatives, strong orders, satisfying both neutrality and the Pareto-condition. We can examine the exact distribution, for which the mean is 18. Without the neutrality condition, but with the Pareto condition, the mean was 16.496. Without the Pareto condition, but with the neutrality condition the mean was 30.95.

Within this class we have failure even of group interjacency. The smallest degree in this class is 6; the highest is 24; degress must be multiples of 6. There exist rules in this class of degree 18, but **none** of degree 12.

C. Anonymity and neutrality

As a last investigation of Pareto rules, we might consider the class of rules with two individuals, three alternatives, strong orders, satisfying all of: anonymity, neutrality and the Pareto condition. But an impossibility theorem prevents analyzing this class: it is empty (see Kelly 1990).

4. Responsiveness

Now let's turn to an explicit responsiveness condition. To express this condition, let's say that profile $u' = (R'_1, \ldots, R'_n)$ *advances* alternative x over profile $u = (R_1, \ldots, R_n)$ if for all y and z different from x, and for all i, $yR'_i z$ if and only if $yR_i z$ and also $xR_i y$ implies $xR'_i y$. In our context, with resoluteness and

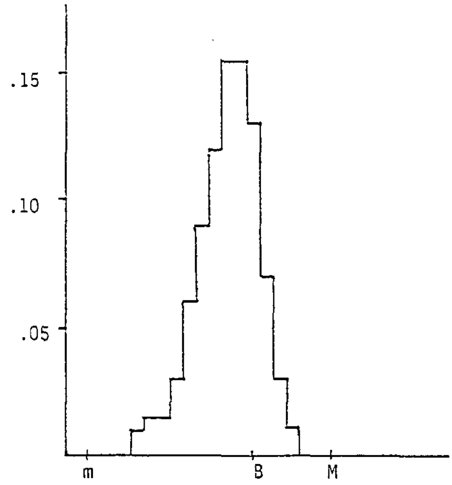

Fig. 12. Degrees of manipulability: Two agents, three alternatives, responsive

universal domain, a social choice rule f satisfies *responsiveness* just when the combination of

1. $f(u)=x$, and
2. u' advances x over u

implies $f(u')=x$. Analysis shows that the minimum degree for the class of responsive rules is 2; the largest degree I have found is 21 (but I do not have a proof that $M=21$). Conditional interjacency holds. I do not have a very efficient program for sampling from the class of responsive rules and I can not yet prove that the sampling procedure is uniform. The inefficiency results in a much smaller sample here: just 1000 rules. The results are presented in Fig. 12. With a sample mean of 13.36, the bunching up near the maximum is not so dramatic and rules are much more spread out. The Borda rule of Fig. 1 satisfies responsiveness. This is the first class where the sample mean is less than the degree (14) of the Borda rule.

A. Neutrality

If we supplement responsiveness with neutrality, we can obtain an exact distribution of manipulability degrees. With two individuals, three alternatives and strong orders, there are only six rules that are both neutral and responsive. Two of these rules have manipulability degree 6 and four have degree 18; the mean is exactly 14. Unsurprisingly, our lowest means arrive with responsiveness.

B. Anonymity and neutrality

Finally, if we supplement responsiveness and neutrality with anonymity, an impossibility theorem strikes again (see Kelly 1990). There do not exist any resolute, neutral, anonymous and responsive social choice rules defined on all profiles of strong orders for three alternatives and two individuals.

References

Black D (1958) Theory of committees and elections. Cambridge University Press, Cambridge

Blair D (1981) On the ubiquity of strategic voting opportunities. Int Econ Rev 22: 649–655

Chamberlin JR (1985) An investigation into the relative manipulability of four voting systems. Behav Sci 30: 195–203

Feldman A (1980) Welfare economics and social choice theory. Martinus Nijhoff, Boston

Gibbard A (1973) Manipulation of voting schemes: a general result. Econometrica 41: 587–602

Kelly JS (1988) Conjectures and unsolved problems: 4. Minimal manipulability and local strategy-proofness. Soc Choice Welfare 5: 81–85

Kelly JS (1989) Conjectures and unsolved problems: 8. Interjacency. Soc Choice Welfare 6: 331–335

Kelly JS (1990) Impossibility results with resoluteness. Econ Lett 34: 15–19

Kemeny J (1959) Mathematics without numbers. Daedalus 88: 577–591

Lacroix SF (1800) Eloge historique de Jean-Charles Borda. Paris

Lepelley D, Mbih B (1987) The proportion of coalitionally unstable situations under the plurality rule. Econ Lett 24: 311–315

Luce RD, Raiffa H (1957) Games and decision: introduction and critical survey. Wiley, New York

Mascart J (1919) La vie et les travaux du Chevalier Jean-Charles de Borda (1733–1799): épisodes de la vie scientifique au XVIII[e] siècle. Lyon

Muller E, Satterthwaite MA (1977) The equivalence of strong positive association and strategy-proofness. J Econ Theory 14: 412–418

Nitzan SI (1985) The vulnerability of point-voting schemes to preference variation and strategic manipulation. Publ Choice 47: 349–370

Nurmi H (1987) Comparing voting systems. Reidel, Dordrecht

Pattanaik PK (1978) Strategy and group choice. North-Holland, Amsterdam

Riker WH (1982) Liberalism against populism. W.H. Freeman, San Francisco

Saari DG (1987) Susceptibility to manipulation. Publ Choice 64: 21–41

Satterthwaite MA (1975) Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J Econ Theory 10: 187–217